Non-Intrusive Electrical Load Monitoring, a Technique for Reduced-Cost Load Research and Energy Management

Leslie K. Norford and Richard D. Tabors, Massachusetts Institute of Technology and Tabors Caramanis and Associates
J. Greg Byrd, Philadelphia Electric Company

Electric utilities have relied on submetering to obtain detailed and accurate information about appliance energy use and peak demand in the residential sector. The cost of submetering and the need to work on the customer side of the revenue meter have limited its use and encouraged the development of a less expensive and less intrusive alternative. A non-intrusive load monitor (NILM) establishes a viable methodology for sampling the voltage and current at the revenue meter at 1 second intervals, detecting changes in real and reactive power, and clustering the changes to aid in identification of individual appliances and disaggregation of the total load. The NILM as tested is designed to detect appliances that are either running at fixed power or turned off, although it can identify the dominant mode of operation of those appliances that operate in several states. This paper describes the results, recently analyzed, of a field test of the NILM in four houses that were submetered for comparison. The ability of the NILM to detect appliances known to be in the test houses, energy use as measured by the NILM and submeters, and potential improvements to the methodology are presented.

Commercial building loads present challenges that were not observed in houses. Many motors, even those running at fixed speed, are subject to variable load and draw variable power. Electrical noise can confound the appliance detector build into the residential NILM. Control signals issued by energy management systems offer valuable information that can be coupled with electrical monitoring. This paper presents results of current work to monitor commercial-building HVAC equipment at a single point. A non-linear filter capable of removing electrical noise is described. Test results from a large office-laboratory facility are summarized. Smoothly varying power levels as measured at a central point were in some cases clearly assigned to a particular device and in other cases required additional information. Applications of the NILM in detecting faulty HVAC equipment performance and optimizing the control of multiple fans, pumps and chillers are assessed.

Introduction

Non-intrusive electric load monitoring provides information about the operation of major electrical equipment by sampling electric power at a single point rather than by submetering. The technique was developed and applied first to residential buildings by researchers who observed that usage patterns of major residential appliances could be detected by reviewing one-second samples of real and reactive power (Hart 1991a, 1991b). This paper presents a first look at recent results from a field test of the non-intrusive load monitor (NILM) as applied to the residential sector, followed by a report of ongoing research that focuses on the commercial sector.

Residential Non-Intrusive Load Monitor

The residential non-intrusive load monitor (Res-NILM) was developed to be a viable alternative to the traditional method of gathering energy use data for residential appliances. Instead of individually metering such appliances as a refrigerator, clothes dryer or hot water heater, the Res-NILM monitors the entire appliance stock from outside the house via a single meter, data recorder and software system. The test of the Res-NILM involved the monitoring of four residences in the service territories of two electric utilities in the Northeast US with the Res-NILM, combined with traditional submetering of specific appliances inside the residences. We briefly describe the Res-NILM system, our analysis techniques, test results and potential improvements.

How the Meter Works

The Res-NILM hardware records changes in total power at the electrical service meter located outside of the home. A microprocessor-based edge detector identifies changes in both real and reactive power (positive and negative) on each leg of the circuit (120V) entering the house, or on both legs (240V). These step changes are due to appliances changing states. The Res-NILM software that we tested was written to find appliances that can be in only two states--on or off. There is, however, a large group of appliances that draw electrical power at more than two discrete levels; these machines will be labeled are multi-state machines (MSMs); both two-state machines and multi-state devices are finite-state machines (FSMs), a category that excludes devices with continuously variable power. MSMs normally contain two or more separate electrical components. Examples of MSM devices are refrigerators, which include a compressor and a defroster, and dishwashers, which have a motor and a heater. Additionally, there are devices such as air conditioners that have different "on" states (low, medium, or high settings). The software developed to identify two-state devices frequently identified the dominant components of MSMs; further work, described in Hart (1991a), has yielded identification algorithms for MSMs.

The Res-NILM meter records those step changes in power that occur within a residence. Because each appliance in a house has specific electrical characteristics, the positive and negative power changes can be associated with individual appliances. Figure 1 illustrates the overall situation by showing hypothetical changes in real power over a 40-minute period.

The prototype Res-NILM system consists, logically, of five steps: measurement; event detection, clustering of events, time-series matching of related on and off events, and appliance identification. The Res-NILM meter installed at each field site measures real and reactive power at one-second intervals. It defines steady power as three one-second-average data points (each a quartet of real and reactive powers on the two branches of residential electrical service) that are constant within an empirically determined tolerance. A start-up transient produces a change that is filtered out until a new, steady power, again defined by three data points, has been established. This efficient, one-pass, edge-detection algorithm was designed to ignore power overshoots associated with motor start-ups and to yield a step change in power. Changes in real and reactive power associated with events are normalized to correct for varying line voltage by scaling the power to what it would have been if the voltage were the nominal 120 volts. This step produces more consistent changes in power for each appliance.

Even with normalization, each appliance exhibits some variation in measured power every time it is switched on. The Res-NILM's processing software groups events into clusters which have approximately the same values for real and reactive power on each of the two residential electrical service legs. Thus one would expect to see all the positive transitions of a hot water heater in one cluster and all the negative changes in another. Each "on" cluster is matched with its corresponding "off" cluster in part by looking at the magnitude of the power (the norm of a four-dimensional vector) and number of events in each cluster. Clustering of data is a type of syntactic pattern recognition that identifies underlying structure in the data (magnitude of real and reactive power in this case) without comparing data to a model via a numerical goodness-of-fit calculation.

Next, the events in a matched pair of clusters are converted to a time series. With perfect data, this step is trivial and produces a sequence in which turn-on and turn-off events alternate. In practice, there are anomalies, where the Res-NILM fails to detect an event and the time series shows sequential on or off events. This problem can arise when appliances have multiple states or when two or more appliances turn on or off at nearly the same time, such that the required steady-state condition for electrical power at the service entrance is not reestablished until after both appliances are operating. The Res-NILM records simultaneous events and their date-time stamp, and uses these events to fill in the anomalies, a useful error-correction scheme.

In order to identify appliances by name, a matching process has been implemented which uses a set of heuristic decision rules that operate on the Res-NILM output (statistics about the matched clusters and the events that make up the cluster). In addition to real and imaginary power levels, these rules relate to voltage level, load type (resistive or inductive), cycle type (continuous or daily) and generic appliance electrical characteristics. These rules, which are currently applied manually, have evolved during the course of the project and work well for the appliances studied. These rules will be further developed as more tests are performed and further experience is gained with the Res-NILM.

The strongest signature characteristic in appliance identification is voltage, 110/120 V (single leg) and 220/240 V (two legs). Dishwashers and refrigerators are 110/120 V appliances, but dryers and heat pumps are 220/240 V

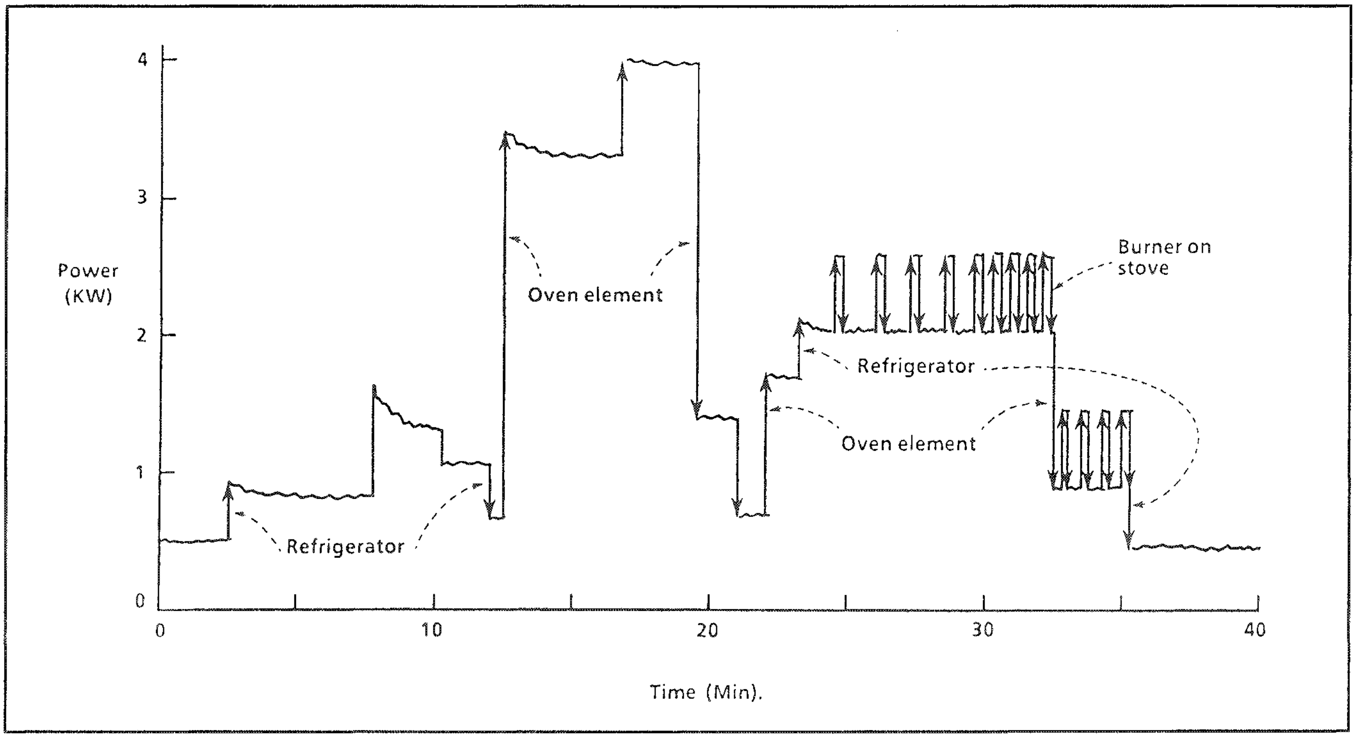

Figure 1. Measured Residential Electrical Power vs. Time, Showing Step Changes Due to Individual Appliance Events (from Hart 1991b)

appliances. Therefore, if a cluster had an average real and reactive power similar to a refrigerator but operates at 240 V, it would not be a refrigerator.

The second signature characteristic is reactive power, present for motors but absent for such devices as resistive heaters. The majority of the 220/240 V loads are resistive, and the majority of 110/120 V are inductive. Knowing that the load is 220/240 and inductive isolates an appliance as a compressor motor for a heat pump or central air conditioner. These two signature characteristics have been combined into a generic appliance electrical characteristic data base which provides typical real and reactive power levels for individual appliance types. Examples are shown in Table 1. The data provided are based on appliance literature; where no literature information was available, Res-NILM data were used.

The third signature characteristic used is cycle type, described as two distinct patterns: *continuous cycling* loads, such as refrigerators, whose on/off pattern is relatively constant throughout the day; and *daily cycling* loads, such as ranges, whose on/off pattern is relatively constant from one day to the next. Certain appliances also exhibit very specific behavior that helps in identifying them. For example, a burner on an electric stove is characterized by the heating element being on for less that 20 seconds and off for less that 30 seconds. Dryers exhibit this behavior but have on-times of about one minute and off-times of about two minutes. Dishwashers and clothes washers also exhibit certain patterns. These patterns are less distinct than the dryer because both the dishwasher and clothes washer are more complex finite state machines than the dryer.

Results

Testing of the Res-NILM system occurred at four sites in the service territories of two electric utilities in the Northeast U.S. Both Res-NILM data for the entire house and traditional meter (parallel) data for specific appliances were collected during the fall, winter and spring of 1987 and 1988. Tabors Caramanis (1992) provides a complete description of the field test.

The sequential on-off time series data formed from matched clusters were compared with submetered data, which were recorded as 5 or 15-minute averages. This comparison revealed that the Res-NILM's error-correction routine did not resolve all of the anomalies due to two successive on or off events. We adopted the method of inserting a missing event at the appropriate time. The time of the event was based on the normal on-off cycle. We did not attempt to determine whether the Res-NILM had

Table 1. Generic Appliance Electrical Characteristics

Appliance	Voltage level	Real Power Watts	Reactive Power Vars
Refrigerator	110/120	100-350	100-550
Pump (sump and pool)	110/120	1200	1200-1100
Dishwasher	110/120	1000	15
Clothes Washer	110/120	500-1000	200-700
Resitive Space Heater	220/240	100-2000	0
Hot Water Heater	220/240	2200-5400	0-40
Stove Burner	220/240	1000-2500	0-40
Dryer	220/240	2500-5500	0-500

found such an event and failed to cluster it, perhaps because of poorly chosen, user-defined parameters that control the clustering process; instead, we simply fabricated an event.

Those anomalies corrected by the Res-NILM are easily seen in the time-series data because the inserted data points are simultaneous events, two or more appliances changing states at the same time, that have larger magnitudes. Thus, the real and reactive power of the simultaneous event are the sum of the appliance transitions and are often very different from the rest of the events that make up the matched cluster pair. We assumed that the resolution of the anomaly by the simultaneous event is appropriate, even though there are cases in which the chosen event is clearly wrong. To preserve the modeled two-state nature of the appliance , the real and reactive components of the simultaneous event were replaced by the average for the matched cluster pair.

Refrigerators are multi-state machines with a compressor which exhibits roughly equal real and reactive power and a defroster which is an electrical resistance heater. Figure 2 provides an excellent example of the match between the Res-NILM and parallel data for the refrigerator at site 4. The large spikes in the parallel data are almost certainly from the defrost unit. Except for the defrost cycle, the Res-NILM correctly identified all on-off cycles shown in the parallel data. The Res-NILM system monitored the refrigerator less well at the three other sites; at two sites, the refrigerators were not found without adjusting one of the clustering parameters.

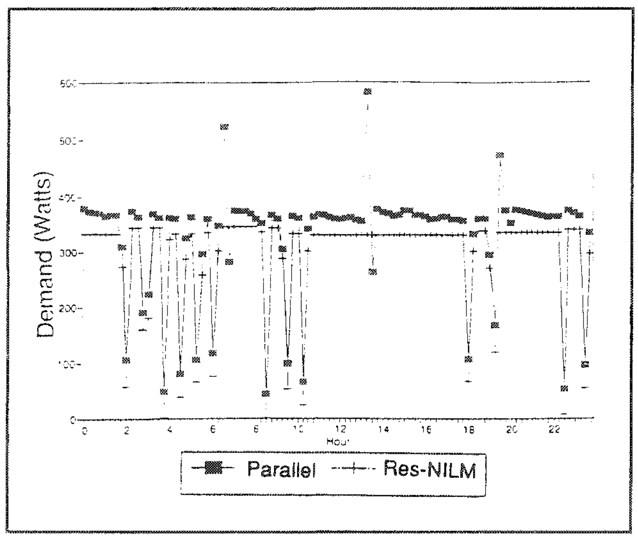

Figure 2. Electrical Power for a Refrigerator as a Function of Time, as Measured by Both the Res-NILM and by an Electrical Meter Attached to the Refrigerator

Electric dryers were identified at all four sites by the Res-NILM. Electric dryers are multi-state devices consisting of heating elements and an electric motor. The current two-state algorithm identified the heating elements and not the motor. The heating elements cycle on and off in order to maintain temperature, a pattern that cannot be seen in Figure 3 because the data are displayed in 15 minute intervals to correspond with the submetered data. Res-NILM and submetered data show that the dryer normally operates for 75 minutes. The software missed a turn-off event during December 20 and a turn-on event on

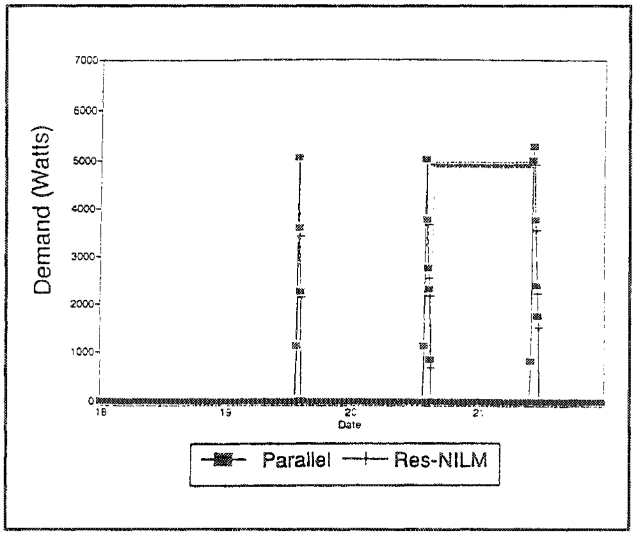

Figure 3. Electrical Power for a Clothes Dryer as a Function of Time, as Measured by Both the Res-NILM and by an Electrical Meter Attached to the Dryer

December 21. Thus, the figure shows the dryer operating from late on December 20 through midday on December 21, which is clearly incorrect. By applying a heuristic that applies the normal heater on-off times of 2 minutes when the dryer is operating, such an obvious error can be corrected, greatly improving Res-NILM's accuracy.

Water heaters, the classic two-state appliance, were accurately monitored at all four sites. The single heating coil operates when the water temperature falls below a certain threshold. Site 4 provides a representative example of how well the software performed. The matched cluster data for a typical day show a perfect match with the parallel data all nine times the water heater operated. The unit usually operated for only about 12 minutes followed by an off-period that ranged from 60 to 240 minutes. The average real power for this water heater was about 2720 watts and the reactive component was about 67 vars. Another indication of how well the software worked in this case was the fact out of 117 pairs of events there were only 2 anomalies.

Electric ranges, on the other hand, were not monitored well with the Res-NILM. The burner cycle times are short relative to the Res-NILM's one-second sample rate and three-second interval to establish steady power, and suggest that the Res-NILM's sampling rate should be increased. If more than one burner element is on, the rapid transitions can be added or subtracted together which complicates the picture. In addition, two burners can generate spurious anomalies by an on-on or off-off sequence. Despite these problems, the Res-NILM identified elements of the burners for the stoves from two sites. Elements of the range at site 4 are believed to have been identified, but the range was not metered at the site. No range was found at site 20.

Dishwashers were identified at the two sites where they were present. The electric dishwasher is a multi-state appliance. It contains a pump that is operating through part of the cycle and a heating element that operates at a different level of demand. The demand identified by the Res-NILM at site 13 was 740 watts while the parallel data showed 850 watts. From the Res-NILM data, the reactive component averaged 460 vars. This combination indicates that either the pump or a combination of the pump and heating elements was detected. While the Res-NILM poorly monitored the site 20 dishwasher, it showed an interesting pattern of on-times and off-times when the dishwasher was operating. The pattern was not perfect which may have been due to the many anomalies and simultaneous events. When the appliance was started, there was a 1-2 minute operating period followed by a 1-2 minute period when the dishwasher was off, followed in turn by a 10 minute operating period and a 10 minute period of no power demand. Usually, there was a final operating period of 7-20 minutes. This pattern is not exhibited by the dishwasher operations incorrectly identified by the Res-NILM. Algorithms that detect such patterns would improve the Res-NILM monitoring accuracy.

Pumps were identified by the Res-NILM at sites 18 and 20. Even though the water pump at site 18 was used often, the Res-NILM captured most of the events. The pump usually operates for about one minute and almost always less than three minutes. The Res-NILM data showed a few operating periods that were more than 30 minutes long. We believed these to be in error, and modified them to be within the normal range of on times (one to three minutes). This hypothesis was confirmed by the parallel data.

Summary

These limited results show that the Res-NILM technique is a viable alternative to traditional metering. The Res-NILM found all of the 25 metered appliances listed in Table 2, which summarizes the submetered and Res-NILM data comparisons. Under the "operations" column in Table 2 are the differences in the number and timing of appliance operations as determined by the Res-NILM and the submetering. This includes times where the Res-NILM did not detect the appliance operation as shown by the parallel data and the times the Res-NILM indicated

Table 2. Comparison of Res-NILM and Parallel Monitoring

Appliance	Site	Operations: Parallel	Operations: Res-NILM Mistakes	Energy (kWh): Parallel	Energy (kWh): Res-NILM	Absolute Error kWh
Refrigerator	4	187	12	109.3	97.5	14.8
	20	8	10	82.6	62.6	21.5
Clothes Dryer	4	7	0	30.8	28.2	3
	18	16	8	--	--	--
	13	8	4	18.9	16.2	17.1
	20	27	5	87.6	84.6	45.4
Hot Water Heater	4	108	0	109	110.9	5.9
	18/1	130/2	13/2	46.9	39.1	21.3
	18/3	77	1	38.3	39.8	17.9
	13	44	4	99.8	94.1	16.4
Electric Range	18	17	5	--	--	--
	13	15	5	--	--	--
Dishwasher	20	32	34	23	23	21.6
	13	3	0	1.8	2.0	0.5
Microwave Oven	18	48	3	2.1	2.2	1.0
	20	23	20	--	--	--
Washing Machine	20	28	4	3.9	3.7	1.5
Water Pumps	18	99	9	3.2	3.7	3.2
	20	108	6	0.16	0.16	0.17
Electric Heaters	13	3/4	3	4.1	4.0	0.5
	18	16	2	12.5	19/5	9/5

/1 120V unit /4 Believe submeter to be in error during part of monitoring period
/2 approximate /5 Includes long on-period not shown by parallel data
/3 220V unit

appliance operation which was not supported by the parallel data. In those cases where the parallel and Res-NILM data are sufficiently close, a comparison of total energy consumed is also useful. However, such a comparison masks combinations of positive and negative differences between the two. To more accurately describe Res-NILM deviations from the parallel data, the "absolute error column in Table 2 gives the absolute value of the difference between the Res-NILM and the parallel data for each 5- or 15-minute period and summed over the analysis period. These period-by-period differences are always greater than or equal to the difference between the total energy determined from the parallel and Res-NILM data. This absolute error also captures imperfect matches between the Res-NILM and parallel data, including differences in appliance start times.

The Res-NILM monitored the appliances listed in Table 2 on a spectrum from excellent for the hot water heaters to poor for the electric range. This indicates the need for further research, not only in identifying multistate machines but also in automatic analysis of time series information to improve identification and resolve anomalies. The comparison between the submetered data and the Res-NILM data also indicates that further field tests should be conducted, not only to broaden the appliance database, but to assist in the inclusion of time series information in the Res-NILM method. Ideally, the parallel data would be recorded at shorter intervals (perhaps every minute) than in the past study, and would include the reactive power component as well.

Commercial Non-Intrusive Load Monitor

For commercial buildings, research now underway seeks to extend the Res-NILM concept in three directions: identification of equipment on the basis of information obtained during start-up transients (and not simply changes in steady-state power); analysis of changes in power level in the context of commercial buildings with an energy management system that turns equipment on and off, making it unnecessary to cluster events as an aid in identification; and analysis of the frequency content of the current drawn by induction motors, as an aid in diagnosing rotor faults. A more detailed description of this research is presented in Norford and Mabey (1992).

Identifying Equipment Based on Start-up Electrical Transients

Because the Res-NILM cannot distinguish different types of electrical equipment with identical steady-state power, it cannot be used in buildings where equipment electrical signals have been corrected to make steady-state power appear to be very similar to purely resistive loads. Electronic power supplies in computers and peripherals, electronic ballasts for gas-discharge lamps, and adjustable-speed motor drives are sources not only of reactive power but also of harmonic distortion that can be suppressed by modifications to the electronics. With these modifications, a bank of fluorescent lamps may, in steady state, look no different from a computer or a space heater. A major thrust of Com-NILM research is the development of pattern recognition tools that are not fooled by stead-state powers that make a load appear as a pure resistor.

Measurements of voltage and current during a turn-on transient (and not, as with the Res-NILM, before and after the transient) reveal patterns characteristic of classes of equipment. Figure 4 shows electrical current for a twin-tube fluorescent light fixture. The lamps show a period of low current during which the mercury in the tubes is vaporized and (not indicated in the figure) current is substantially out of phase with voltage, followed by a larger current (nearly in phase with the voltage) after the arc is struck. In contrast, current for a motor starts at a high value to accelerate the rotor and drops down to a steady-state level required to overcome constant-speed losses.

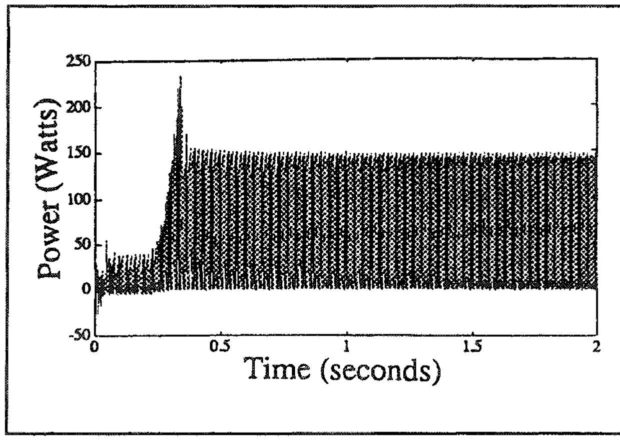

Figure 4. Electrical Power for a Fluorescent Lamp Fixture, During the Start-Up Transient

Numerical pattern discriminators can compare the start-up transient of an unknown device with patterns known to be manifested by different types of equipment. Because the duration of start-up transients may depend on the size of the equipment (large motors have longer time constants than small motors), this comparison must take place at multiple time resolutions. A goodness-of-fit indicator, such as a Euclidean norm or a vector inner product, that exceeds a threshold serves to establish a match and therefore identify a device. In essence, the information gained from the start-up transient adds a dimension to the clusters established for the Res-NILM, thereby distinguishing clusters of similar power but distinct start-up behavior. The Com-NILM being developed is therefore a hybrid pattern recognition device, combining both numerical discriminant and syntactic pattern recognition approaches.

By working with the details of the transients, the Com-NILM may be able to distinguish devices for which the start-up transients overlap. Here the key concept is that there are only brief periods during the transient when power is actually changing, separated by periods of nearly steady power. In the time between these information-rich periods, the Com-NILM is ready to detect the start of another device.

The Com-NILM as Part of a Building Automation System

A second commercial-building application we have investigated, in field and laboratory tests, posits the presence of an automatic control system responsible for the start-up and shut-down of major HVAC equipment. By limiting, at this stage of our research, the role of the NILM to that of support for an automation system, we have access to important information, namely the identity of a device and the precise time at which the control system initiates a start-up or shut-down signal. Therefore, changes in power can be assigned to the device receiving the control signals, and the strength of the association can be tested statistically. There is no longer a need to make an association on the basis of clustering procedures that must be at once sufficiently inclusive to lower the risk of ignoring step changes caused by the appliance in question, and sufficiently exclusive to discriminate appliances that have comparable values for clustering parameters. The primary advantage of relying on start-stop signals rather than clusters for device identification is enhanced fault detection: a change in power indicative of a fault might fall outside a cluster, but, with little or no ambiguity, can be related temporally to a control signal.

Field Test. The field test consisted of taking one-second average measurements of real power from the electrical distribution panel that provides power to the HVAC equipment in two campus buildings. The equipment includes two identical 500-ton centrifugal chillers and associated chilled water (50 hp) and condenser water (40 hp) pumps, two large supply fans with adjustable-speed drives (125 and 100 hp), and a number of smaller pumps and fans; maximum total power for the system was about 1000 kilowatts. We employed the data to both detect equipment on-off transitions and to pinpoint two types of equipment-operation faults.

Detection of changes in power involved two changes from the Res-NILM metering concept. First, we increased the width of the detection window, to reduce the uncertainty associated with measured changes in power. With the measured standard deviation of about 5 kW and a sample size of 10, an HVAC device drawing more than 5 kW can be detected at the 5% level of significance.

The second change involved an exploration of filtering algorithms more time consuming than the one-pass edge detector used in the Res-NILM. The Res-NILM filter is sensitive to power spikes, whether or not the spikes are associated with a change in power level. That is, a spike above a specified tolerance will cause the filter to reset, compute a new average power after the spike, and assign the before-and-after difference in power to a cluster. The filter works well if spikes are indeed connected to changes in power level, as is the case with induction motors that draw a large current during the start-up transient. If, however, the spike is an event that is not associated with a start-up transient, it will cause an unnecessary computational burden and force the analyst to subsequently discriminate against a cluster that indicates essentially no change in power. A loose tolerance that passes spikes will not solve the problem if the spikes are as large as power levels of devices that the Com-NILM is meant to detect. This problem can be reduced when equipment control signals are available, by turning off the event-detection algorithm at other times, but such an approach requires an additional layer of logic and cannot eliminate spikes that occur during a start-up or shut-down event.

The ideal filter for the Com-NILM would remove both isolated spikes and spikes associated with motor dynamics. We found that a median filter satisfactorily met both objectives. The median filter is one of a class of nonlinear filters that incorporate a sorting function (Karl et al. 1990). It sorts the data points in its window and preserves the median before sliding over one point and repeating the process. This filter removes impulses, which appear in the window as one or two points of very large or small power that have little or no effect on the median. The sort can be time-consuming but as compensation offers the significant benefit of requiring only a single parameter, the width of the window, which can be determined readily by the minimum temporal feature of interest. We used a window of 11 points.

We took data for both a pump and a fan. Figure 5 illustrates four on and off transitions for a 50 hp condenser-water pump, which was tested at a time when the other condenser pump was in constant operation. The increase in system power when the pump is running is large relative to the 5 kW standard deviation and periods of pump operation can be visually discerned. However, the figure shows that the pump motor is more complicated than a simple two-state device with a motor start-up power surge, for several reasons:

1. After start-up, there is a long period of slowly decreasing power, during which fluid pressures in the water loop are approaching equilibrium. The pump just started must reach a balance with the pump already running, at which time each pump is drawing less power than either would if running alone.

2. After the test pump is turned off, the remaining pump, still in operation, is subjected to a larger load and its power increases. There is only a narrow window in time--about 5 seconds--during which to establish the power change of the test pump at shutdown; a larger averaging period will be affected by the increased power of the second pump. We have observed similar behavior with the two chillers.

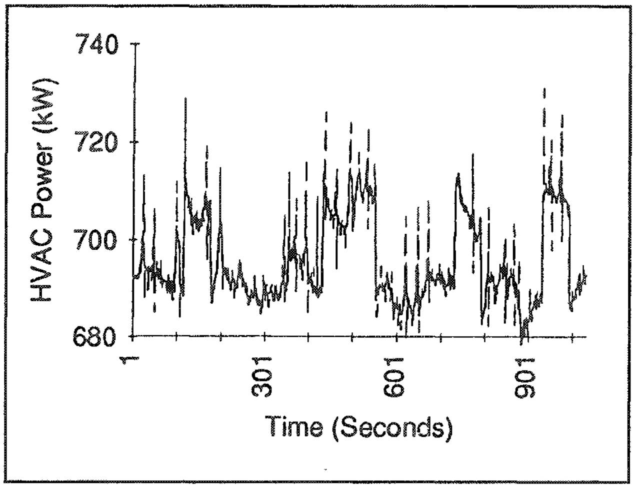

Figure 5. Changes in Electrical Power at the HVAC Service Entrance of a Building, Due to Switching a Condenser-Water Pump On and Off. Noise in the raw data (dashed) is removed by the median filter (solid).

3. There are periodic spikes as large as 20 kW magnitude, which may be due to variable-speed-drive fan controllers responding to set-point changes. These spikes motivated our investigation of the median filter and demanded its application. Figure 5 shows the extent to which the median filter removed unwanted information and Table 3 compares the median filter against the Res-NILM filter. Included in Table 3 are results from a filter we devised that calculates the integral of the power after a increase above a threshold. This approach also rejects spikes and retains step changes. The median filter performs as well, based on a comparison of standard deviations and requires only a single parameter, making it much less sensitive to characteristics of the data set under analysis. Further tests, to be performed, will include submetering to provide a comparison of the calculated changes in power during start-up and shut-down events.

The fan motors are controlled by variable-speed drives (VSDs) that include a soft-start feature to eliminate start-up power surges. The start-up power is a ramp that successive applications of the median filter, with a widening window, would convert to a sharp edge, simplifying the task of estimating start-up power but masking detail in the power signal subsequent to start-up.

Applications in HVAC Fault Detection. We describe particular types of faults that our non-intrusive metering found and assess additional applications that rely on more rapid sampling.

High Controller Gains Figure 6 shows large power oscillations with a peak-to-peak amplitude, about 150 kW, that exceeds the rated power of any device except the chillers. This simple, heuristic identification relies on a list of equipment and rated powers. Oscillations started when the total power dropped and stopped when the power rose or, as shown in Figure 6, when a piece of equipment identified from the control signals as one of the two chillers was turned off. The data led us to conclude that oscillations occurred at times when the chillers were lightly loaded and were due to poorly tuned controller gains. A Fast-Fourier-Transform of the data showed a strong spike at the frequency associated with the four-minute sample time used by the chilled-water temperature controller.

Switching Multiple Chillers Optimal control strategies have been developed by Braun et al. (1989) and detection of power deviations from optimal conditions has been explored by Pape et al. (1991). The amount of information required for optimization has deterred its acceptance by industry and building owners. A NILM offers a lower cost, somewhat less informative, but still powerful approach by providing a basis for identifying what is clearly not correct, even if it is not possible to establish how to achieve what is optimal. The earlier discussion of chiller oscillations is one example of this approach.

A second example concerns switching among multiple chillers and chilled water pumps. Braun showed that optimal switch points can be defined as producing no discontinuity in system coefficient of performance (COP). That is, if a second chiller is turned on at too low a cooling load, COP will drop and power will rise. Power will drop if it is turned on at too high a cooling load. The same argument can be made for pump switching. The data in Figure 6 show that the second chiller was turned off after the optimal cooling load switch point had been passed. Mean power dropped by about 100 kW, indicating that the combined power for both chillers exceeded the optimum by that amount for some period of time leading up to shutting down the second chiller. This information, if detected by a NILM, can guide plant operators toward more efficient plant operation.

A type of chiller staging not observed experimentally in this study concerns when to turn on the first chiller. At issue is how large a cooling load can be met by the ventilation system bringing in 100% outdoor air, and how much fan power is required. Traditional practice has maintained the supply air temperature at a fixed value, forcing the chiller on when the outside temperature approaches this set point (with a small decrement due to temperature rise across the fan). Fan power will therefore stay the same immediately after the chiller is turned on

Table 3. Changes in Electrical Power Due to Turning a Pump on and off, as Determined by Application of Three Different Filters

	Median (11 point)		Integral (10 kW)		Res-NILM (5 kW)		Median+Res-NILM	
Trial	On	Off	On	Off	On	Off	On	Off
1	18.2	-17.6	24.2	-19.8	18.4	-16.5	15.6	-14.5
2	20.9	-18.2	23.5	-16.7	25.1	-25.9	20.6	-20.5
3	18.8	-17.3	22.4	-19.6	15.2	-18.1	16.0	-18.8
4	21.7	-20.3	25.0	-18.7	29.6	-24.7	24.3	-19.0
								-18.2
Average	19.9	-18.3	23.8	-18.7	22.1	-21.3	19.1	
Std Dev	1.7	1.4	1.1	1.4	6.5	4.7	4.1	2.6

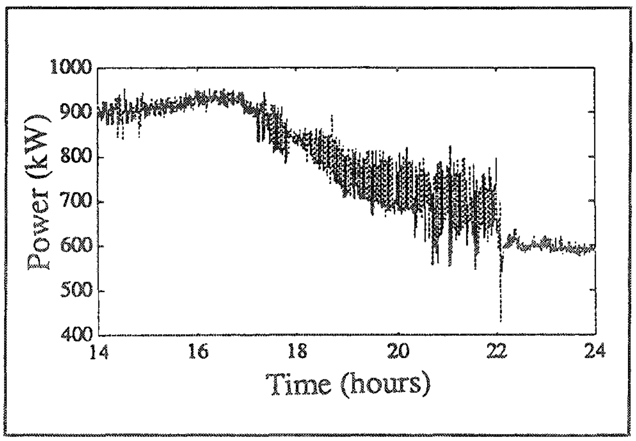

Figure 6. Electric Power at the Building HVAC Service Entrance. A poorly tuned chiller controller operating under low-load conditions caused the large power oscillations.

and the chiller will be running at relatively low load. Alternatively, the supply air temperature could be allowed to float upward, with the chiller turned on when the fan is running at maximum load or (less likely) when the increase in fan power exceeds the power drawn by the chiller. The latter case is exactly the same as the problem of staging the second chiller and the NILM will tell whether there has been a telltale net increase in power.

Motor Faults

Almost all of the HVAC facility managers interviewed as part of this project indicated that the failure of motors used to drive pumps, fans, or compressors was a significant problem. Motor failure was described as unanticipated and disruptive. Tavner and Penman (1987) state that monitoring the electrical current drawn by the stator of an induction motor is one method of assessing the condition of motors, in order to prevent unanticipated failures. Such monitoring detects rotor defects, which are of particular interest because rotor imperfections can deteriorate under the large electrical, mechanical and thermal stresses present in the rotor. In frequency space, rotor-bar defects yield stator-current harmonics displaced from the fundamental.

The technique has also been used to detect damaged gear boxes and rotor eccentricity due to bearing wear. Liu (1990) examined, via theory and experiment, stator-current indications of eccentricity in the airgap separating stator and rotor. She distinguished two types of eccentricities: fixed deviations from perfect concentricity, due to mounting the rotor shaft slightly off center; and dynamic deviations, due to either an out-of-round rotor or bearing damage, that yield an airgap which, at any point, will vary as the rotor turns. Liu found that static eccentricity drops the current at the fundamental frequency, dynamic eccentricity yields a single frequency sideband, lower than the fundamental, and that static and dynamic eccentricities in combination produce a series of harmonics. Because the harmonic amplitudes were difficult to predict, Liu concluded that detection of motor faults requires analysis of how harmonics change over time.

We have analyzed electrical current harmonic spectra in a laboratory setting, for combinations of three motors. Two were fractional horsepower single phase motors and the third a larger three-phase motor. The shaft of one of the single-phase motors was connected via a belt under tension to another pulley, subjecting the rotor to a torque that would tend to create eccentricity. The prominent sidebands located at about 7.5 Hz above and below the fundamental, shown in Figure 7, were caused by the belt and disappeared when the belt was disconnected. These sidebands were also visible when both of the single-phase motors were operating concurrently. However, when the third, larger motor was then turned on, its harmonic spectrum obscured the sidebands associated with the smaller motor. As a next step, we will take current spectra in our test building and determine whether sidebands can be observed when motors are run individually (as in many cases they could be, during periods of time when the remaining HVAC equipment was shut down) or simultaneously.

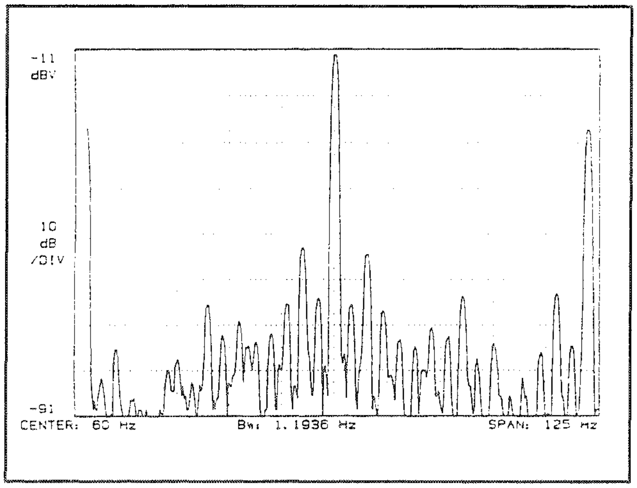

Figure 7. Motor Harmonic Spectrum, with a Belt Attached to the Shaft of the Motor. The belt, under tension, creates a torque on the shaft that causes rotor eccentricity, as revealed in the sidebands to the fundamental signal.

Successful current analysis by the Com-NILM requires sampling rates higher than that used in the Res-NILM but comparable to that required to analyze start-up transients. Of more concern is the question of how to relate changes over time in the magnitude of the electrical sidebands to failure probabilities for motors, a topic now under investigation.

Conclusion

Non-intrusive electrical load monitoring is a technique for sampling and analyzing building electrical signals at central locations, with the goal of replacing expensive submeters. For houses, field tests showed that the meter worked well when installed at the electrical service entrance, detecting nearly all of the major appliances. A comparison with submetered data taken during the tests shows that improvements can be made in detecting appliances with multiple operating conditions and in recognizing typical operating patterns of common appliances.

Work on a meter for commercial buildings has identified and, to an extent, tested, several enhancements to the meter and several new applications. Sampling at rates significantly higher than the 1 Hz frequency used in the residential meter permits more positive device detection on the basis of characteristic start-up transients and also supports frequency analysis of motor currents that can identify faults in the motor's rotor. Slow-speed (1 Hz) sampling, combined with filtering to remove unwanted electrical noise, has proved capable of detecting HVAC equipment, eliminating the need for individual current transducers to inform building automation systems that equipment is in operation, and has also proved to be an effective tool in finding equipment that is switched on or off at sub-optimal times or equipment with poorly tuned controllers.

Further tests of the residential non-intrusive meter are needed to better quantify its accuracy. The commercial metering concepts have been tested with off-line analysis of data and require both further work and codification into an on-line device.

Acknowledgments

The authors gratefully acknowledge financial support and technical advice provided by the Electric Power Research Institute and Johnson Controls. The residential non-intrusive load monitor was designed by George Hart, now at Columbia University. Kurt Levens, Ruel Little, Andrea Kendrick and Alice Yates assisted with the research on the commercial-building meter. Jim Kirtley, Steve Leeb and George Verghese have performed much of the work on the commercial-building meter (notably the analysis of start-up transients) and have provided guidance for the efforts described in this paper.

References

Braun, J. E.; Klein, S. A.; Mitchell, J. W.; and Beckman, W. A. 1989. "Applications of Optimal Control to Chilled Water Systems without Storage." *ASHRAE Transactions* Vol. 95, Pt.1.

Hart, G.W. 1991a. "Nonintrusive Appliance Load Monitoring." Department of Electrical Engineering and Center for Telecommunications Research, Columbia University.

Hart, G.W. 1991b. "Advances in Nonintrusive Appliance Load Monitoring." *Proceedings of the Information and Automation Conference*, Washington, D.C. Electric Power Research Institute.

Liu, S. 1990. "Detection of Airgap Eccentricities in Squirrel Cage Induction Motors Using Stator Current Measurements." M.S. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA.

Karl, W. C.; Leeb, S. B.; Jones, L. A.; Kirtley, J. L.; and Verghese, G. C. 1990. "Applications of a Class of Nonlinear Filters to Problems in Power Electronics." *Record of the IEEE Power Electronics Specialists Conference.*

Norford, L. K. and Mabey, N. 1992. "Non-Intrusive Electrical Load Monitoring in Commercial Buildings." *Proceedings of the Eighth Symposium on Improving Building Systems in Hot and Humid Climates*, Texas A&M University, Dallas, Texas.

Pape, F. L. F.; Mitchell, J. W.; and Beckman, W. A. 1991. "Optimal control and fault detection in heating, ventilating, and air-conditioning systems." *ASHRAE Transactions*, Vol. 97, Pt.1.

Tabors Caramanis and Associates. 1992. "Nonintrusive Load Monitoring: Draft Final Report." Electric Power Research Institute, Palo Alto, CA.

Tavner, P. J. and Penman, J. 1987. *Condition Monitoring of Electrical Machines*. Research Studies Press, Exeter, England.